

CLAIMS:

- 1 1. A method for managing an incoming data message at a host node in a switched
2 fabric, comprising:
3 determining whether if there are pre-post receive buffers specified for a client upon
4 registration by the client;
5 if there are pre-post receive buffers specified for the client, posting client specified
receive buffers at management queue pairs (QPs) to receive the incoming data message; and
 if no pre-post receive buffers are specified for the client, posting a default number of
receive buffers at the management queue pairs (QPs) to receive the incoming data message.
- 2 2. The method as claimed in claim 1, further comprising:
3 continuing to monitor and receive incoming data messages, and determining if the
number of posted receive buffers falls below a threshold value; and
4 if the number of posted receive buffers falls below the threshold value, posting additional
5 receive buffers to receive the incoming data messages.
- 1 3. The method as claimed in claim 2, further comprising:
2 if the number of posted receive buffers exceeds an upper threshold value, removing a
3 designated number of receive buffers posted to receive the additional incoming data messages so

1 as to conserve resources.

1 4. The method as claimed in claim 3, further comprising:
2 monitoring a receive buffer usage of the client based on the number of incoming data
3 messages received for the client are received; and
4 increasing the number of receive buffers posted on behalf of the client to receive the
5 number of incoming data messages intended for the client.

5. The method as claimed in claim 2, wherein the default value of number of receive
buffers is set by a fabric administrator based on operating conditions of the switched fabric,
including a fabric size and a traffic pattern.

6. The method as claimed in claim 2, wherein the threshold value is set by a fabric
administrator based on operating conditions of the switched fabric, including a number of local
clients registered at the host node in the switched fabric.

1 7. The method as claimed in claim 2, wherein the management queue pairs (QP) are
2 unreliable datagram (UD) queue pairs, including QP0 managed by an agent of subnet services,
3 known as Subnet Management Agent (SMA), and QP1 managed by the agent of general services,
4 known as General Services Agent (GSA), in accordance with the "*InfiniBand™ Architecture*

1 *Specification".*

1 8. A host system comprising:
2 at least one channel adapter (CA) including one or more ports to support data transfers,
3 via a subnet; and
4 an access module including a General Services Agent (GSA) and a Subnet Management
5 Agent (SMA) to enable one or more entities to send and receive data messages of management
6 services on the host system, via the subnet, including to determine an optimal number of receive
7 buffers to post at management queue pairs (QPs) so as to receive an incoming data message from
8 the subnet.

9 9. The host system as claimed in claim 8, wherein one of the General Services Agent
10 (GSA) and the Subnet Management Agent (SMA) is configured to:

11 determine whether if there are pre-post receive buffers specified for a client upon
12 registration by the client;

13 if there are pre-post receive buffers specified for the client, post client specified receive
14 buffers at the management queue pairs (QPs) to receive the incoming data message; and

15 if no pre-post receive buffers are specified for the client, post a default number of receive
16 buffers at the management queue pairs (QPs) to receive the incoming data message.

1 10. The host system as claimed in claim 9, wherein one of the General Services Agent
2 (GSA) and the Subnet Management Agent (SMA) is further configured to:

3 monitor and receive incoming data messages, and determine if the number of posted
4 receive buffers falls below a threshold value; and

5 if the number of posted receive buffers falls below the threshold value, post additional
6 receive buffers to receive the incoming data messages.

11. The host system as claimed in claim 10, wherein one of the General Services
Agent (GSA) and the Subnet Management Agent (SMA) is further configured to remove a
designated number of receive buffers posted to receive the additional incoming data messages so
as to conserve resources, if the number of posted receive buffers exceeds an upper threshold
value.

12. The host system as claimed in claim 11, wherein one of the General Services
Agent (GSA) and the Subnet Management Agent (SMA) is further configured to monitor a
receive buffer usage of the client based on the number of incoming data messages received for
the client are received, and increase the number of receive buffers posted on behalf of the client
to receive the number of incoming data messages intended for the client.

13. The host system as claimed in claim 10, wherein the default value of number of

1 receive buffers is set by a fabric administrator based on operating conditions of the subnet,
2 including a subnet size and a traffic pattern.

1 14. The host system as claimed in claim 10, wherein the threshold value is set by a
2 fabric administrator based on operating conditions of the subnet, including a number of local
3 clients registered at the host system in the subnet.

1 15. The host system as claimed in claim 10, wherein the management queue pairs
2 (QP) are unreliable datagram (UD) queue pairs, including QP0 managed by the Subnet
3 Management Agent (SMA), and QP1 managed by the General Services Agent (GSA), in
4 accordance with the "*InfiniBandTM Architecture Specification*".

1 16. The host system as claimed in claim 10, wherein the management services include
2 a subnet administration service which provides data path information to reach fabric-attached
3 devices; a communication management service which provides the means to set up and manage
4 communications between queue pairs (QP); a performance management service which specifies a
5 set of facilities for examining various performance characteristics of the subnet; a device
6 management service which specifies the means for determining the type and location of various
7 types of fabric-attached devices; a device configuration service which assigns fabric-attached
8 devices to the host system; a baseboard management service which allows management of the

1 fabric-attached devices; and a network protocol service which specifies mechanisms to support
2 transport of Simple Network Management Protocol "SNMP" operations through the subnet.

1 17. A computer readable medium comprising instructions that, when executed by a
2 host system in a switched fabric including end nodes and switches interconnected via links, cause
3 the host system to determine an optimal number of receive buffers to post at management queue
4 pairs (QPs) by performing the steps of:

determining whether if there are pre-post receive buffers specified for a client upon
registration by the client;

if there are pre-post receive buffers specified for the client, posting client specified
receive buffers at management queue pairs (QPs) to receive the incoming data message; and

if no pre-post receive buffers are specified for the client, posting a default number of
receive buffers at the management queue pairs (QPs) to receive the incoming data message.

1 18. The computer readable medium as claimed in claim 17, further comprising
2 instructions that, when executed by the host system, cause the host system to further perform the
3 steps of:

4 continuing to monitor and receive incoming data messages, and determining if the
5 number of posted receive buffers falls below a threshold value; and

6 if the number of posted receive buffers falls below the threshold value, posting additional

1 receive buffers to receive the incoming data messages.

1 19. The computer readable medium as claimed in claim 18, further comprising
2 instructions that, when executed by the host system, cause the host system to remove a
3 designated number of receive buffers posted to receive the additional incoming data messages so
4 as to conserve resources, if the number of posted receive buffers exceeds an upper threshold
5 value.

20. The computer readable medium as claimed in claim 18, further comprising
instructions that, when executed by the host system, cause the host system to monitor a receive
buffer usage of the client based on the number of incoming data messages received for the client
are received, and increase the number of receive buffers posted on behalf of the client to receive
the number of incoming data messages intended for the client.

1 21. The computer readable medium as claimed in claim 18, wherein the default value
2 of number of receive buffers is set by a fabric administrator based on operating conditions of the
3 switched fabric, including a fabric size and a traffic pattern.

1 22. The computer readable medium as claimed in claim 18, wherein the threshold
2 value is set by a fabric administrator based on operating conditions of the switched fabric,

1 including a number of local clients registered at the host node in the switched fabric.

1 23. The computer readable medium as claimed in claim 18, wherein the management
2 queue pairs (QP) are unreliable datagram (UD) queue pairs, including QP0 managed by an agent
3 of subnet services, known as Subnet Management Agent (SMA), and QP1 managed by the agent
4 of general services, known as General Services Agent (GSA), in accordance with the
5 "*InfiniBandTM Architecture Specification*".